

FY2005 Accomplishments

ESnet

ESnet Network Accomplishments

William E. Johnston (wej@es.net), ESnet Manager, Lawrence Berkeley National Laboratory

Summary

The Energy Sciences Network (ESnet) enables the large-scale research of the Office of Science (SC)--it is driven by science requirements, as explored in SC-sponsored workshops, for increased network performance and new network services such as guaranteed bandwidth. In FY2005 bandwidth to many ESnet sites was substantially increased. The completion of the San Francisco Bay Area Metropolitan Area Network provides the Bay Area laboratories with 20-30 Gb/s bandwidth as well as multiple connectivity for high reliability. The ESnet On-demand Secure Circuits and Advance Reservation System (OSCARS) project has demonstrated the feasibility of providing end-users with guaranteed bandwidth between sites.

ESnet is driven by the science requirements of SC. Two workshops¹ systematically developed these requirements as they relate to networking and middleware by examining a set of major science disciplines for which the process of science related to computing and communication must change over the next decade in order to make significant progress. ESnet has developed a strategy for meeting the science requirements and has started implementation. This involves:

- 1) Increasing the bandwidth and reliability of network access available to DOE researchers by building the San Francisco Bay Area MAN (BAMAN) and the west coast segments of a second national core – the Science Data Network (SDN).
- 2) Providing guaranteed bandwidth as a service by building a system to automatically schedule and implement virtual circuits traversing ESnet and the other R&E networks.
- 3) Improving the ability of scientists to access network measurement data for all network segments end-to-end that are critical to their science by participating in an international collaborative measurements effort.

1. Increasing network bandwidth and reliability

The completion of BAMAN enabled redundant 10Gbps production IP (Internet) access, and additional 10-20 Gbps access for large science data flows, to the five DOE sites in the SF Bay Area – SLAC, LBNL, NERSC, JGI, LLNL and SNL/CA. The hubs for both the production IP core and SDN core national networks are also included in order to provide site connections to two independent backbones. The BAMAN, with its redundant ring architecture and connections to two backbones, replaced singly connected local loops ranging from 600 Mbps to 2.5 Gbps. The access rates and reliability are now approaching what is required of the network for the large-scale science of the future.

The completion of the west coast segments of the SDN (figure 1) is the start of the second national backbone that will provide failover for the IP core network and dedicated bandwidth for large science data flows. The new segments increase the access to ESnet from 155 Mbps to 10Gbps for General Atomics and PNNL and provide for 10Gbps connection to Pacific Wave and to the high-speed (10Gbps) traffic exchanges for international R&E networks.

¹ “High-Performance Networks for High Impact Science” www.es.net/#research .

2. Implementing guaranteed bandwidth service

With the Large Hadron Collider (LHC) expected to come online in 2007, demand for guaranteed high-bandwidth connectivity for huge data transfers is becoming urgent. The ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARS) is

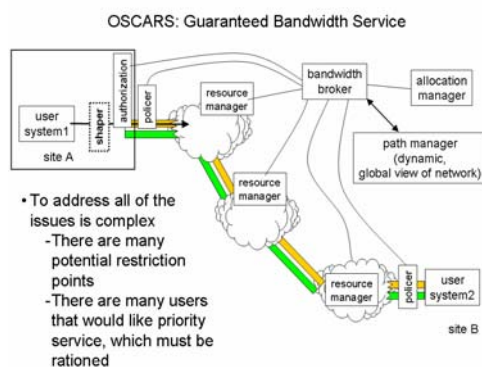


Figure 1 Showing components of the OSCARS system.

developing and deploying a prototype service that enables on-demand provisioning of guaranteed bandwidth, secure virtual circuits within the ESnet production network. The service is being developed in collaboration with Internet2/Abilene and the US R&E network community, and with DANTE/GÉANT and the European R&E network community.

In the past year, ESnet has successfully demonstrated that an end-user can provision circuits within ESnet (through a simple Web-interface), and effectively obtain the required bandwidth. To date we have created 21 accounts for beta test users, collaborators and developers, and have processed more than 100 reservation requests.

3. Enhancing network measurements

Network measurement is critical for the success of widely distributed applications that move large amounts of data for debugging and tuning the application and the network.

The ESnet measurements project has joined with the Internet2 PIPES project and the GÉANT Joint Research Activity 1 project in a collaborative effort to develop a generalized network measurement framework. These are the networks that are in the paths of almost all science data flows. The effort is critical for effective use of high-speed networks by high performance distributed applications, particularly international collaboration. The first version of the framework design is completed and an early prototype has been developed and deployed in ESnet, GÉANT and Internet2 supports sharing network link

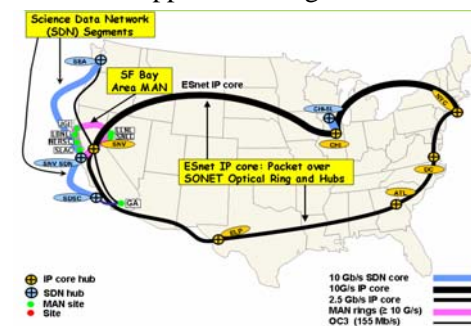


Figure 2 Highlighting the Bay Area MAN and the west coast segment of SDN

capacity and utilization data. The instances deployed are already being used by the Enabling Grids for E-science project for developing Grid software for network operators. The framework uses standardized schemas to facilitate interoperability with other network measurement research projects.

For further information on this subject contact:
Mary Anne Scott, Program Manager
Mathematical, Information, and Computational Sciences Division
Office of Advanced Scientific Computing Research
Phone: (301) 903-6368
scott@er.doe.gov

ESnet PKI Accomplishments

Michael Helm, ESnet - helm@es.net

Summary

ESnet provides the DOEGrids Certificate Authority (CA) for Public Key Infrastructure based identity authentication, supporting a number of Office of Science programs. The DOEGrids CA also supports DOE scientific collaborations with other federal agency research programs, as well as international collaborations such as the CERN LHC. In support of these collaborative, cross-agency and cross-boundary activities ESnet is helping to organize regional and global certificate authority Policy Management Authorities that operate under policies defined by the science community.

DOEGrids Certificate Authority

The DOEGrids Certificate Authority signs X.509 identity (authentication) certificates for people and Grid services (computer processes and computer hosts) involved in collaborative science. The CA policies are determined and governed by the science community, and DOEGrids is one of largest issuers of certificates to Grid users and hosts.

The basic purpose of the CA is to extend into cyberspace the traditional web of trust (policies) that enables large-scale, global scientific collaborations. This service is provided for twelve different “Registration Authorities (RA)” that roughly coincide with Grid Virtual Organizations or collaborations. They include Grid efforts at six DOE national laboratories or sites; several scientific collaborations. Virtual organizations like iVDGL and

PPDG are highly dependent on this service and makes heavy use to provide both “people” and “service” certificates. Several sites, most visibly Fermi Lab (see Figure 1), make heavy use of DOEGrids for host and Grid service certificates.

The DOEGrids, initiated in 2001, has seen rapid growth in the number of host and service certificates issued and substantial growth in people certificates issued (see figure 2) this year.




Figure 2 - Certificates Issued

In response to its customers, DOEGrids is continually streamlining and improving certification processes. Interface scripting and automation is one aspect of streamlining, and is particularly effective for issuing host and service certificates.




Figure 1 - DOEGrids VO Breakdown (“others” are certificate renewals)

Other CA and PKI Activities

ESnet operates a self-signed root CA which signs subordinate certificate authorities, one of which is the DOEGrids CA. ESnet also operates two custom certificate authorities for the DOE NERSC supercomputer center and for the National Fusion Collaboratory, which are separately developing new PKI and user management services that fall outside of what DOEGrids CA policy allows. The ESnet root CA has also signed a subordinate CA operated by NERSC to support one-time-passwords and Kerberos integration development.

An experimental certificate revocation service (OCSP) is provided to test on-demand, network-based certificate validity testing. This service will advance to production status in the next year as requirements and innovations from European physics collaborations mature. DOEGrids oversight is provided by the DOEGrids Policy Management Authority (PMA), which is composed of representatives of all of the served science collaborations, and which meets on a quarterly basis.

Federation

In order to promote compatibility in the science community the project has been working with European Grid providers and CA operators, and has maintained a membership in the European PMA governing CA practices. This organization is now known as the [EUGridPMA](#) and is closely affiliated with a European Grid management body, the [EGEE](#). There are a great many Grid projects in the world for which DOE has affiliations that are outside of the immediate scope of DOEGrids. To insure interoperability among science Grids in the US and elsewhere, the project is aggressively promoting and sponsoring two policy boards.

The Americas Grid PMA ([TAGPMA](#)) will provide roughly the same level of services in the Americas (US, Canada, and Latin America) that EUGridPMA provides for European Grid CAs. TAGPMA builds on that experience and will focus on supporting innovation and interoperability. Several very large US-based Grid consortia are also looking to TAGPMA to provide standards, CA evaluations, and CA registries in order to minimize the burden of this work on their programs. TAGPMA is operated by CANARIE, a Canadian R&D organization. The newly formed International Grid Trust Federation ([IGTF](#)) is operated by ESnet. The IGTF interfaces the various Grid regional PMAs that are developing, such as EUGridPMA and TAGPMA. Standards issues and innovations originating in each regional PMA are coordinated through the IGTF. IGTF coordinates its meetings and activities with the Global Grid Forum meetings, providing a publishing vehicle for standards such as CA profiles, operational best practices, and other documents. It is anticipated that DOEGrids will offer a great many services through IGTF that will flatten out the hierarchy of boards and organizations that have developed; these might include unified CA repositories, directories, revocation information, and trouble ticket / problem dispatch. However, the requirements and implementation details are the subject of discussion and development in the coming year.

For further information on this subject contact:

Mary Anne Scott, Program Manager
Mathematical, Information, and
Computational Sciences Division
Office of Advanced Scientific Computing
Research
Phone: (301) 903-6368
scott@er.doe.gov

ESnet Audio, Video, and Data Collaboration Services

William E. Johnston (wej@es.net), ESnet Manager, Lawrence Berkeley National Laboratory

Summary

ESnet Audio, Video, and Data Collaboration Services (AVD) are used by over 1000 scientists and researchers worldwide in order to see, hear and exchange information with remotely located collaborators. These services increase productivity and reduce costs by reducing travel and telephony expenses. The ESnet AVD project is committed to providing the latest voice, video and data collaboration technology and applications to its customers to further increase their productivity and reduce costs.

ESnet AVD collaboration services support voice, video, and data collaboration technology which provides DOE Office of Science researchers and their collaborators the ability to meet and exchange information remotely as easily as if they were in the same location. With the potential of tens of thousands of customers, ESnet AVD technology has evolved from a telephony-based manual operation to a primarily IP-based automated technology infrastructure that is easily scalable to support many thousands of users.

At the present time, the ESnet AVD service has over 1000 registered users worldwide supporting such science initiatives as High Energy Physics projects (ATLAS, D0, CDF, ILC, CMS, ZEUS, OSG, DOSAR), Magnetic Fusion projects (Alcator, C-Mod) and others. The standards based IP-video conferencing service (H.323) provides scientists over 5000 port hours per month (with monthly and seasonal variations), the audio conferencing (telephony) provides over 2000 port hours per month, and data conferencing provides 100 to 200 port hours per month. (Each person in conference uses one port.)

New technology is being tested to provide voice over IP (VoIP) technology for AVD customers. IP-based voice meetings reduce the dependency on telephony and increase savings by using ESnet's IP network for telephony meetings.

The AVD services leverage ESnet to provide DOE scientists with very substantial cost savings. For example, Fermi Lab use of IP-based video conferencing reduced their telephony costs from over \$12,000 per month to less than \$100 per month. Variations of this scenario have occurred at other DOE labs.

Architecture

ESnet AVD Collaboration service consists of the following components:

1. Web-based registration (<http://www.ecs.es.net>) where potential customers register themselves, if required, their equipment, and obtain help.
2. IP videoconferencing (H.323) provides 180 ports of video capability and a centralized video "switch" called a gatekeeper. Customers are given ESnet assigned video numbers to use and are free to meet anytime in a purely "ad-hoc" manner.

3. Audio and data conferencing is still telephony based and is a scheduled and reserved service. A customer schedules a meeting at a web site. Each meeting participant receives email notification of the meeting time. The audio bridge consists of 144 telephone ports for users. Data conferencing supports thousands. One of the important aspects of the service for science collaboration is a globally accessible, centralized meeting scheduling service

In January, 2004, ADV changed the video collaboration service from one based on telephony to one based primarily on the ESnet IP network. Since then tools have been developed to track the usage for all ADV services.

Once a year, the ADV Collaboration project holds a Workshop so users can provide direction and advice. The Workshop 2004 is documented at

<http://hpcrd.lbl.gov/ESnetCollab>

The trends for all services show steady or increasing usage over time. ESnet expects increased usage as more people become aware of the AVD Collaboration Services.

For further information on this subject contact:

Mary Anne Scott, Program Manager
Mathematical, Information, and
Computational Sciences Division
Office of Advanced Scientific Computing
Research
Phone: (301) 903-6368
scott@er.doe.gov

ESnet Network Progress and New Directions

William E. Johnston (wej@es.net), ESnet Manager, Lawrence Berkeley National Laboratory

Summary

The Energy Sciences Network (ESnet) is driven by the requirements of the science Program Offices in DOE's Office of Science (SC). To that end, ESnet provides a high-bandwidth network connecting forty-two DOE sites to each other and to collaborators worldwide. Each year, the Office of Science facilities are used by more than 18,000 researchers from DOE Labs, universities, other government agencies, and private industry. As DOE's large scale science continues to move to a distributed international model, ESnet is providing the innovation and expertise to meet its networking needs. This entails not only increasing end-to-end bandwidth, but becoming actively involved with other domestic and international research and education (R&E) networks in developing and deploying cutting edge operational technologies.

Overview

ESnet is evolving to meet the needs of DOE science as identified in an August 2002, DOE Office of Science-sponsored workshop¹. This evolution is tracking the results of a follow-on workshop, held in June 2003², which constructed a roadmap describing the network and related middleware services necessary to meet the ambitious networking needs of DOE's large-scale science.

A New Architecture

A new ESnet architecture and a new implementation strategy have been developed, and the next-generation network is being incrementally deployed to increase the bandwidth, services, reliability and cost effectiveness of the network. The elements of the architecture include multiple, independent, optical channel-based, national

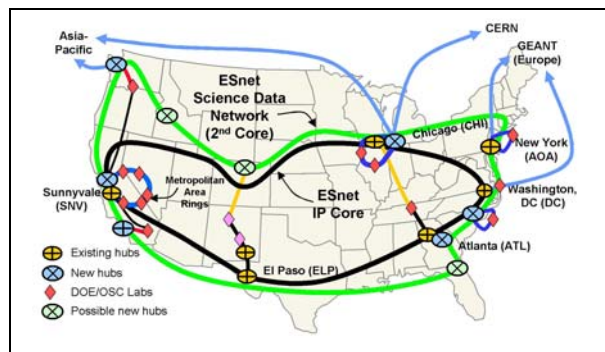
cores (each at 10-40 Gb/s) that independently connect to Metropolitan Area Network (MAN) rings, together with independent paths to the major R&E networks of Europe and Japan by connecting to all of the available peering points.

The MAN rings provide redundant paths and on-demand high bandwidth point-to-point circuits for DOE Labs. The multiple cores connect to the MAN rings in different locations to ensure that the failure of a core node could not isolate the MAN. This is illustrated in the figure using the current ESnet IP core and the Science Data Network (SDN) core which is built from National Lambda Rail (NLR)³ optical channels. The first two segments of the second core – 10 Gb/s circuits from San Diego to Sunnyvale, CA to Seattle have been put into service.

¹ "High-Performance Networks for High Impact Science." Report of the August, 2002, Workshop Conducted by the Office of Advanced Scientific Computing Research of the U.S. Department of Energy Office of Science

² "DOE Science Networking Challenge: Roadmap to 2008." Report of the June, 2003, DOE Science Networking Workshop. Both Workshop reports are available at <http://www.es.net/#research>.

³ An advanced services network of the US research and education community.



Another aspect of the new architecture is high-speed peering with the US university community via the Internet2/Abilene network. US universities are an important component of DOE science and require state-of-the-art access to the DOE laboratories served by ESnet.

Implementation Strategy

The implementation involves building the network by taking advantage of the evolution of the telecom milieu – that is, using non-traditional sources of fiber, collaborations with existing R&E network confederations for lower cost transport, and vendor-neutral interconnect points for more easily achieving financial competition for the “last mile” tail circuits to ESnet sites.

Replacing the current point-to-point tail circuits with MAN optical rings is providing high-speed, high-quality production IP service, at least one backup path from DOE labs to ESnet hubs, scalable bandwidth options from sites to the ESnet core, and point-to-point provisioned high-speed circuits as an ESnet service. The newly completed SF Bay Area MAN connects five DOE sites to both national core networks and is the first ESnet MAN.

Involvement with the Networking R&D Community

A clear mandate from the Roadmap Workshop was that ESnet should be more

closely involved with the network R&D community, both to assist that community and to more rapidly transition new technology into ESnet. To facilitate this, the new implementation strategy includes multiple interconnection points with NLR based test beds and UltraScienceNet – DOE’s network R&D testbed.

ESnet has been very active in collaborating with the R&D community, the European R&E network, DANTE/ GÉANT, and the domestic R&E network, Internet2/Abilene, in the areas of applied network research that are directly applicable to creating the seamless end-to-end paradigm required for science. Specifically, the OSCARS⁴ project that dynamically creates end-to-end private virtual networks is a collaboration in order to ensure an interoperable, inter-domain approach that will allow scientists to run the specialized protocols needed to move vast quantities of data between the various networks. ESnet’s collaboration with the R&D community on perfSONAR⁵, an inter-domain monitoring framework, is assisting in assuring that the end-to-end paths are functioning correctly. ESnet is also sharing its expertise by participating in the DHS sponsored Secure Routing Workshops, whose purpose is to secure the fundamental reachability protocols on which the entire Internet is dependent.

For further information on this subject contact:

Mary Anne Scott, Program Manager
Mathematical, Information, and
Computational Sciences Division
Office of Advanced Scientific Computing
Research

⁴ OSCARS- On-demand Secure Circuits and Advance Reservation System <http://www.es.net/oscars/index.html>

⁵ perfSonar- Performance Service Oriented Network monitoring Architecture

ESnet RADIUS Authentication Fabric: Solving the authentication delivery problem

Michael Helm, ESnet - helm@es.net

Summary

ESnet has prototyped a RADIUS Authentication Fabric, to link together and federate existing authentication services in DOE laboratories and collaborating institutions. The RAF was developed to support various one-time password initiatives under study at various DOE laboratories in early 2004, but can also be applied to many other large scale interoperability problems, such as WAN wireless roaming.

Secure authentication is one very important aspect of improved cyber security. Strong identity verification involves two factor authentication where one of the factors is typically a hardware cryptographic token (smart card, USB device, stand-alone challenge-response device, etc.). The user provides a system name and user name and the token supports a unique challenge-response for every login.

ESnet RAF Development

The RADIUS authentication fabric supports cross-site operation of crypto token based authentication devices which are typically issued at, and the authentication requests validated only at the user's home institution.

The RAF has a core set of servers operated by ESnet that essentially function as authentication routers. They receive authentication queries using the RADIUS protocol, which is a widely deployed and commercially supported authentication and authorization protocol standardized by the IETF. The ESnet RADIUS servers know how to route these authentication queries to the appropriate destination. This might be another DOE laboratory, or a collaborating university site, or even another RAF hierarchy in operation in Internet2 or in Europe. While RADIUS has some security capabilities of its own and native support for basic authentication types, it is also capable of piggybacking more sophisticated and secure protocols and thus serving as an authentication transport mechanism.

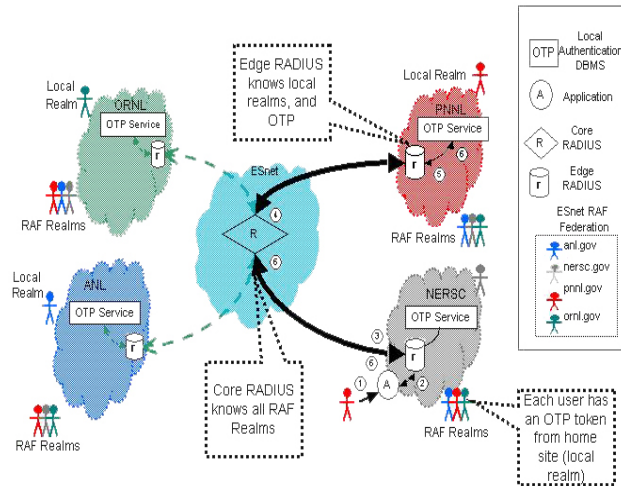
At the request of the ESSC (ESnet's steering committee), ESnet prototyped this architecture

and tested it for federating one-time password initiatives under study at NERSC, ORNL, LBNL, and several other places. The project demonstrated that it could use the RAF to eliminate the need for both sites and individuals to support multiple tokens for their cross-site collaborations, essentially providing a one-time password single-sign-on solution. The report and proposed architecture can be found at the ESnet RAF website - <http://www.es.net/raf>.

ESnet RAF Future

Interest in one-time password solutions remains strong in the DOE community. Many "virtual organizations" or large-scale, cross-site projects have considerable interest in one-time passwords, but are particularly sensitive to the burden of multiple organizations and multiple service providers, and the federation opportunity provided by the RAF is an attractive solution. ESnet will continue to support these efforts in the Fusion and High Energy Physics communities in the coming year.

In the past year the RAF project has developed relationships with a similar project in Europe ([Eduroam](#)), and an early-stage effort in Internet2 (the Federated Wireless NetAuth (FWNA) working group), to support wireless roaming across multiple academic institutions. The current plan is to interconnect the ESnet RAF with FWNA and Eduroam on an experimental basis, and help develop solutions to the many significant scaling and security issues that remain.



For further information on this subject
contact:

Mary Anne Scott, Program Manager
Mathematical, Information, and
Computational Sciences Division
Office of Advanced Scientific Computing
Research

Phone: (301) 903-6368

scott@er.doe.gov